

Computation of winning strategies for μ -Calculus by fixpoint iteration

Christian Neukirchen


TCS Oberseminar · July 11, 2014

Overview

- Short introduction to μ -calculus
- Parity games and strategies
- Strategies for μ -calculus
- Example: mutex
- Implementation and optimization
- Future

Labelled Transition Systems

We consider LTS having a non-empty set of states S , total relations $\xrightarrow{a} \in S \times S$ (for actions $a \in \mathcal{A}$) and propositions $p \in \mathcal{P}$ which hold at a state or not.



μ -calculus: Syntax (from Hofmann and Ruedel 2014)

$\phi ::= X$	
p $\neg p$	
$[a]\phi$	<i>(for all a-transitions)</i>
$\langle a \rangle \phi$	<i>(a-transition exists)</i>
$\phi_1 \wedge \phi_2$ $\phi_1 \vee \phi_2$	
$\mu X.\phi$	<i>(least fixpoint)</i>
$\nu X.\phi$	<i>(greatest fixpoint)</i>

Only propositions can be negated to ensure monotonicity.

μ -calculus: Examples

“q holds everywhere along the path”

$$\nu X.q \wedge [a]X$$

“q holds infinitely often on the path”

$$\nu X.\mu Y.(q \wedge \langle a \rangle X) \vee \langle a \rangle Y$$

“p holds at every even position” (more powerful than CTL*)

$$\nu X.p \wedge \langle a \rangle \langle a \rangle X$$

μ -calculus: Set semantics

$$\text{sem}(X, \eta) = \eta(X)$$

$$\text{sem}(\phi_1 \wedge \phi_2, \eta) = \text{sem}(\phi_1, \eta) \cap \text{sem}(\phi_2, \eta)$$

$$\text{sem}(\phi_1 \vee \phi_2, \eta) = \text{sem}(\phi_1, \eta) \cup \text{sem}(\phi_2, \eta)$$

$$\text{sem}([a]\phi, \eta) = \widetilde{\text{pre}}(\overset{a}{\rightarrow})(\text{sem}(\phi, \eta))$$

$$\text{sem}(\langle a \rangle \phi, \eta) = \text{pre}(\overset{a}{\rightarrow})(\text{sem}(\phi, \eta))$$

$$\text{sem}(\mu X. \phi, \eta) = \text{iter}_X(\phi, \eta, \emptyset)$$

$$\text{sem}(\nu X. \phi, \eta) = \text{iter}_X(\phi, \eta, S)$$

$$s \in \widetilde{\text{pre}}(\overset{a}{\rightarrow})(U) \Leftrightarrow \forall t \in S. s \overset{a}{\rightarrow} t \implies t \in U$$

$$s \in \text{pre}(\overset{a}{\rightarrow})(U) \Leftrightarrow \exists t \in S. s \overset{a}{\rightarrow} t \wedge t \in U$$

$$\text{iter}_X(\phi, \eta, U) = \text{let } U' := \text{sem}(\phi, \eta[X := U]) \text{ in}$$

$$\text{if } U = U' \text{ then } U \text{ else } \text{iter}_X(\phi, \eta, U')$$

Parity games

A *parity game* consists of a disjoint sum of positions $\text{Pos} = \text{Pos}_0 \cup \text{Pos}_1$, a total edge relation $\rightarrow \subseteq \text{Pos} \times \text{Pos}$ and a priority function $\Omega : \text{Pos} \rightarrow \mathbb{N}$.

Moves happen along the edge relation. The destination decides who moves next.

The game is *won* if the largest priority that occurs infinitely often is even, the opponent wins if it is odd.

Strategies for parity games

A *strategy* ρ is a function that tells the player how to move next.

A *positional strategy* only takes the the current position into account.

A position is in a *winning set* W_i if there exists a strategy ρ such that player i wins, starting at a position in W_i .

Theorem 1. Every position p is either in W_0 or W_1 and player i wins positionally from every position in W_i .


Strategies for μ -calculus

We can interpret a μ -calculus formula ϕ as a parity game. Moves can happen along the subformulae (example next slide). The priority of a position depends on the kind of formula and its nesting depth.

A partial winning strategy for μ -calculus is a partial function

$$\begin{array}{ll} \Sigma : \Phi \times S \rightarrow s & \text{(move to state } s \in S) \\ | 1 & \text{(take the left formula)} \\ | 2 & \text{(take the right formula)} \\ | * & \text{(take the only formula)} \end{array}$$

Small strategy example



$$\nu X. \mu Y. (q \wedge \langle a \rangle X) \vee \langle a \rangle Y$$

$$X := \nu X. Y$$

$$Y := \mu Y. (q \wedge \langle a \rangle X) \vee \langle a \rangle Y$$

$$X \ 0 \ \rightarrow \ * \ |$$

$$X \ 1 \ \rightarrow \ * \ |$$

$$Y \ 0 \ \rightarrow \ * \ |$$

$$Y \ 1 \ \rightarrow \ * \ |$$

$$q \ 1 \ \rightarrow \ * \ |$$

$$\langle a \rangle X \ 0 \ \rightarrow \ X \ 1 \ |$$

$$\langle a \rangle X \ 1 \ \rightarrow \ X \ 1 \ |$$

$$\langle a \rangle Y \ 0 \ \rightarrow \ Y \ 1 \ |$$

$$\langle a \rangle Y \ 1 \ \rightarrow \ Y \ 1 \ |$$

$$q \ /\ \ \langle a \rangle X \ 1 \ \rightarrow \ * \ |$$

$$(q \ /\ \ \langle a \rangle X) \ \backslash / \ \langle a \rangle Y \ 0 \ \rightarrow \ \#2 \ |$$

$$(q \ /\ \ \langle a \rangle X) \ \backslash / \ \langle a \rangle Y \ 1 \ \rightarrow \ \#1$$

Updating strategies

Given two winning strategies Σ and Σ' , we can define the partial winning strategy $\Sigma + \Sigma'$ as

$$\begin{aligned}(\Sigma + \Sigma')(\phi, s) = & \text{if } (\phi, s) \in \text{dom}(\Sigma) \\ & \text{then } \Sigma(\phi, s) \\ & \text{else } \Sigma'(\phi, s)\end{aligned}$$

Strategy semantics

$$\text{SEM}(X)_\eta = \{(X, s) \mapsto * \mid s \in \eta(X)\}$$

$$\text{SEM}(p)_\eta = \{(p, s) \mapsto * \mid p \text{ holds at } s\}$$

$$\text{SEM}(\neg p)_\eta = \{(p, s) \mapsto * \mid p \text{ does not hold at } s\}$$

$$\begin{aligned} \text{SEM}(\phi \wedge \psi)_\eta &= \text{SEM}(\phi)_\eta + \text{SEM}(\psi)_\eta \\ &\quad + \{(\phi \wedge \psi, s) \mapsto * \mid (\phi, s) \in \text{dom}(\text{SEM}(\phi)_\eta) \\ &\quad \quad \quad \wedge (\psi, s) \in \text{dom}(\text{SEM}(\psi)_\eta)\} \end{aligned}$$

$$\begin{aligned} \text{SEM}(\phi \vee \psi)_\eta &= \text{SEM}(\phi)_\eta + \text{SEM}(\psi)_\eta \\ &\quad + \{(\phi \vee \psi, s) \mapsto 1 \mid (\phi, s) \in \text{dom}(\text{SEM}(\phi)_\eta)\} \\ &\quad + \{(\phi \vee \psi, s) \mapsto 2 \mid (\psi, s) \in \text{dom}(\text{SEM}(\psi)_\eta)\} \end{aligned}$$

$$\text{SEM}([a]\phi)_\eta = \text{SEM}(\phi)_\eta \\ + \{([a]\phi, s) \mapsto * \mid (\phi, s) \in \text{dom}(\text{SEM}(\phi))_\eta\}$$

$$\text{SEM}(\langle a \rangle \phi)_\eta = \text{SEM}(\phi)_\eta \\ + \{(\langle a \rangle \phi, s) \mapsto s' \mid s \xrightarrow{a} s' \wedge (\phi, s') \in \text{dom}(\text{SEM}(\phi))_\eta\}$$

$$\text{SEM}(\nu X.\phi)_\eta = \text{SEM}(\phi)_{\eta[X:=\text{sem}(\phi, \eta)]}$$

$$\text{SEM}(\mu X.\phi)_\eta = \text{ITER}_X(\phi, \eta, \{\})$$

$$\text{ITER}_X(\phi, \eta, \Sigma) = \text{let } \Sigma' := \text{SEM}(\phi)_{\eta[X:=\text{dom}(\Sigma)]} \text{ in} \\ \text{if } \Sigma = \Sigma' \text{ then } \Sigma \text{ else } \text{ITER}_X(\phi, \eta, \Sigma')$$


Checking strategies

Easy algorithm for checking whether a strategy is correct:

Run the strategy until you find a loop (hit the same formula at the same state again): then check whether the highest priority inside the loop is even (good) or odd (strategy is wrong!).

Can be implemented as simple recursive traversal.

Example: mutual exclusion (from Huth and Ryan 2004)



Safety “Only one process is in its critical section at any time.”

$$\forall Z. \neg(c_1 \wedge c_2) \wedge [a]Z$$

Liveness “Whenever any process requests to enter its critical section, it will eventually be permitted to do so.”

$$\forall Z. (\neg t_1 \vee (\mu X. c_1 \vee [a]X)) \wedge [a]Z$$

Mutex safety: sample run

```
% ./micromu.native huth-fig3.7.lts huth-fig3.7-safety.mu  
Z =nu/1= ~c1 \/ ~c2 /\ ([a]Z)
```

```
### Execution time Mu.sem: 0.000000s  
result: 0 1 2 3 4 5 6 7
```

```
### Execution time Strat.sem: 0.003333s  
...
```

```
verifying for good state 0:  
TRAV Z , 0 , loop-search  
TRAV ~c1 \/ ~c2 /\ ([a]Z) , 0 , loop-search  
TRAV ~c1 \/ ~c2 , 0 , loop-search  
TRAV ~c1 , 0 , loop-search
```

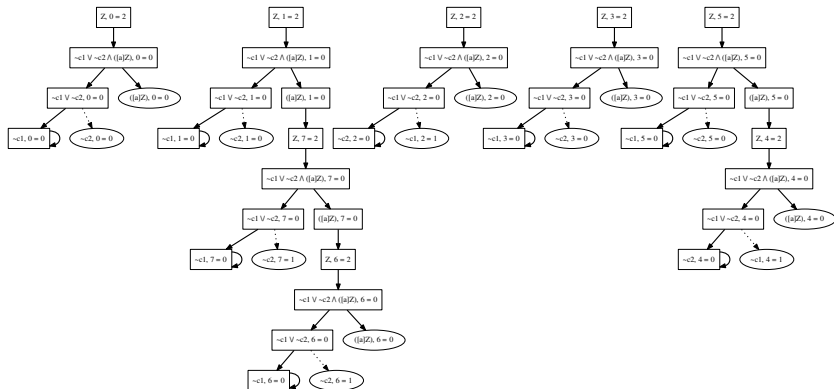


```
TRAV ~c1 , 0 , loop-search
TRAV ~c1 , 0 , maxprio 0 here 0
TRAV ~c1 , 0 , maxprio 0 here 0
TRAV ([a]Z) , 0 , loop-search
state 0: true
...
verifying for good state 5:
TRAV Z , 5 , loop-search
TRAV ~c1 \/ ~c2 /\ ([a]Z) , 5 , loop-search
TRAV ~c1 \/ ~c2 , 5 , loop-search
TRAV ~c1 , 5 , loop-search
TRAV ~c1 , 5 , loop-search
TRAV ~c1 , 5 , maxprio 0 here 0
TRAV ~c1 , 5 , maxprio 0 here 0
TRAV ([a]Z) , 5 , loop-search
TRAV Z , 4 , loop-search
TRAV ~c1 \/ ~c2 /\ ([a]Z) , 4 , loop-search
```

```
TRAV ~c1 \/ ~c2 , 4 , loop-search
TRAV ~c2 , 4 , loop-search
TRAV ~c2 , 4 , loop-search
TRAV ~c2 , 4 , maxprio 0 here 0
TRAV ~c2 , 4 , maxprio 0 here 0
TRAV ([a]Z) , 4 , loop-search
state 5: true
verifying for good state 6:
state 6: known good
verifying for good state 7:
state 7: known good

### Execution time verify: 0.000000s
result:  0 1 2 3 4 5 6 7
verification passed
```

Mutex safety: strategy run



Mutex liveness: sample run

```
% ./micromu.native huth-fig3.7.lts huth-fig3.7-liveness.mu
X =mu/4= c1 \/ ([a]X)
Z =nu/1= ~t1 \/ X /\ ([a]Z)
```

```
### Execution time Mu.sem: 0.000000s
result:
```

The formula holds nowhere: the strategy is empty. The mutex does not guarantee liveness. Why?


Mutex liveness: generating counterexamples

To generate a counter example, we can tell micromu to negate the formula using `-c`:

```
% ./micromu.native -c huth-fig3.7.lts huth-fig3.7-liveness.mu
### Execution time Mu.sem: 0.000000s
### Execution time Strat.sem: 0.003333s
### Execution time gendot: 0.006666s
### Execution time verify: 0.003333s
result:  0 1 2 3 4 5 6 7
verification passed
```



Looking at the strategy, we can find the counterexample.

Mutex liveness: counterexample strategy



The big loop corresponds to
 $0 \rightarrow 1 \rightarrow 3 \rightarrow 7 \rightarrow 1 \rightarrow$
 $3 \rightarrow 7 \rightarrow \dots$

Fixing the mutex



```
% ./micromu.native huth-fig3.8.lts huth-fig3.7-liveness.mu
```

```
...
```

```
result: 0 1 2 3 4 5 6 7 8
```

```
verification passed
```

Implementation

compact: about 1 kLOC OCaml (another 1 kLOC thrown away during development), no external dependencies

quick: “worst-case” exponential example¹ \mathcal{G}_1 takes 0:02:40 and uses 28 MB RAM (down from 3+ hours / 6+ GB...)

simple, recursive algorithms: verifier should be easy to port to proof assistant (Coq)

¹15 nested quantifiers, cf. Friedmann 2009, section 5ff.

Systems of equations

First version actually substituted variables inside μ -formulae: consumes exponential amount of memory with nested formulae.

Rewrite equations as a ordered system of equations:

$$\nu Z.(\neg t_1 \vee (\mu X.c_1 \vee [a]X)) \wedge [a]Z$$

$$Z \stackrel{\nu}{=} (\neg t_1 \vee X) \wedge [a]Z$$

$$X \stackrel{\mu}{=} c_1 \vee [a]X$$

Need order to restore original formula.

Vastly simplifies implementation: makes ν -case trivial, μ -case a lot easier.

Optimizations

- Caching of results for $\nu X.\phi$ case
- Avoiding OCaml polymorphic compare (`compare_val`)
- Using maps for strategies, not association lists (requires careful strategy update)
- Caching of verified states (else easily quadratic runtime)
- Very helpful: `ocamlcp(1)`/`ocamlprof(1)` and `perf(1)`

Future: formal verification

- The checker is meant to be a *certified decision procedure*
- Formally verifying the checker to be correct results in a verified implementation of μ -calculus
- Done so far: definitions of least and greatest fixpoints (on arbitrary sets), specialized version of Knaster-Tarski, μ -calculus set semantics
- To do: serialize strategies into Coq terms
- To do: implement checker for strategies (using finite sets)
- To do: prove checker correct
- To do: extract verified checker?

Questions?

Thank you.

References

- [1] Oliver Friedmann. “An Exponential Lower Bound for the Parity Game Strategy Improvement Algorithm as We Know it”. In: *LICS*. IEEE Computer Society, 2009, pp. 145–156. ISBN: 978-0-7695-3746-7.
- [2] Martin Hofmann and Harald Rueß. “Certification for mu-calculus with winning strategies”. In: *ArXiv e-prints* (Jan. 2014). arXiv: 1401.1693 [cs.LG].
- [3] Michael Huth and Mark Dermot Ryan. *Logic in Computer Science: Modelling and Reasoning About Systems*. 2nd. New York, NY, USA: Cambridge University Press, 2004, pp. I–XIV, 1–427. ISBN: 052154310X.